

UNCLASSIFIED

Defense Technical Information Center Compilation Part Notice

ADP023091

TITLE: Tracker Analysis and Ground Truth Tool Description for the Proceedings of the 2002 Ground Target Modeling and Validation Conference

DISTRIBUTION: Approved for public release, distribution unlimited

This paper is part of the following report:

TITLE: Proceedings of the Ground Target Modeling and Validation Conference [13th] Held in Houghton, MI on 5-8 August 2002

To order the complete compilation report, use: ADA459530

The component part is provided here to allow users access to individually authored sections of proceedings, annals, symposia, etc. However, the component should be considered within the context of the overall compilation report and not as a stand-alone technical report.

The following component part numbers comprise the compilation report:

ADP023075 thru ADP023108

UNCLASSIFIED

Tracker Analysis and Ground Truth Tool Description for the Proceedings of the 2002 Ground Target Modeling and Validation Conference

Carrie Kimbel, Jeff Todd, Mark A. Chambliss, Judson R. Griffin III,
Daniel Konkle, Paul D. Lavalley, Dr. Jay Lightfoot
Dynetics, Inc.
Huntsville, AL 35814 USA

ABSTRACT

The Tracker Analysis and Ground-Truth (TAG) tool was developed to streamline the process of defining a target region for each image in an image sequence in order to calculate signature metrics and perform tracker analysis. Previously, images within a sequence were analyzed individually by outlining the target in each image and calculating the relevant statistics. Likewise, tracker performance was determined by a man-in-the-loop viewing of tracker results overlaid on the image sequence. While effective, these methods did not allow the timely processing of long sequences and limited the calculation of signature metrics to only a few frames of a sequence. Using the TAG tool, a single rectangle is drawn around the target at the beginning of the sequence. The rectangle is shifted and resized using mouse motion and keystrokes as the sequence is played. The rectangle size and location is stored in a file which can be replayed and modified. This file can be used as input to SEMIRS (Software for Extracting Metrics from Infrared Sequences) for generation of signature metrics for each image and as stabilization input to ISAT (Infrared Seeker Algorithm Tool) which generates tracker output. The TAG tool has made it possible to generate, store, and utilize valuable ground truth information for entire sequences in a short period of time. This paper will focus on the utilization, functionality and application of the TAG tool.

INTRODUCTION

The process of analyzing infrared test imagery often requires ground-truth information, or the location and size of a target region. Ground-truth information is required to calculate signature metrics and to evaluate the performance of tracker algorithms. The target location and size define the target pixels that are to be used in metric calculations, and quantify tracker performance by determining the overlap between the target region and the track gate.

Prior to the creation of the TAG Tool, several commercial software packages were used in a piecewise fashion to generate the ground-truth data. First, image sequences were broken into individual image files in order to be imported in Scion Image (a.k.a. NIH Image). Each image was opened separately in Scion Image and an outline was drawn around the target. The pixels inside the outline were set to a value of two and the pixels outside the outline were set to zero. The modified image was then saved as the ground-truth file for that image. When complete, individual ground-truth files existed for every image in the sequence. It should be obvious that this method was very inefficient.

The TAG Tool was developed to streamline the process of generating ground-truth information by processing image sequences as a unit. To generate ground-truth information, a single rectangle is drawn around the target at the beginning of the sequence. Once outlined, ground-truth mode is invoked and the rectangle is resized and repositioned on every frame as the sequence is played. The rectangle can be resized using keystrokes and moved using mouse motion while in this mode. The location and size of the box are recorded for each frame and may be saved in a single text file. In this manner, ground-truth information for an entire sequence may be generated in a reasonable amount of time. For a sequence of about 3000 frames, it takes the typical user about 2 hours, while the manual method would take about 40 hours.

OPENING AN IMAGE SEQUENCE

In order to ground-truth an image sequence, the sequence must first be opened. The TAG Tool will open several formats: ADIF (Air Defense Image Format), NDIF (New Air Defense Image Format), bin

(Raytheon format), or raw. The raw format is especially useful because it is user-defined. A file may be opened using the 'Open ...' or 'Import ...' command under the 'File' menu as shown in *Figure 1*, or using

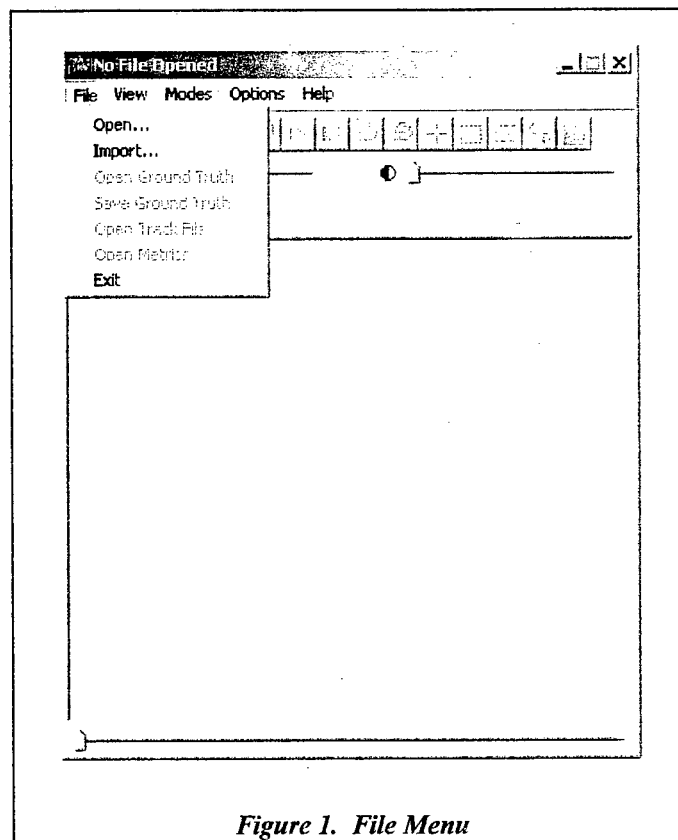


Figure 1. File Menu

toolbar button with the file folder icon. The 'Import File' command should only be used for raw image types. It will automatically bring up the raw image properties dialog box once a file is chosen. Files may also be opened by simply dragging and dropping them into the image display area. In this case, the file type will be determined using the file's extension. If the extension does not match a known format, it will be opened as a raw format. Once a file is opened, the name of the file will be displayed in the title bar.

PLAYBACK CONTROLS

The toolbar has seven VCR-like playback control buttons that control the playback of the image sequence. From left to right they are: play backward, rewind, step backward, step forward, go to frame, play forward, and stop. The 'go to frame' button will bring up a dialog asking for the destination frame number. If a specific frame number is not needed, but only a general location within the sequence, the slider bar underneath the image may be used to quickly advance to a general location within the sequence.

Keystrokes may also be used to play through the sequence. The stand-alone arrow keys are used for this and are applied as follows: up arrow = play forward, down arrow = stop, right arrow = step forward, left arrow = step backward. Play advance may be set to 1, 2, 5, 10, or 20 frames per step under the 'Options' menu. This allows for a quicker viewing of the entire sequence.

The current frame number is displayed directly above the right hand side of the image. By default the frames are displayed counting down, starting with the total number of frames. This order may be reversed by using the 'Options' pulldown menu and selecting 'Count Up' under 'Image Numbers.' Two other pieces of information are also displayed directly above the image: the value in the center is the current zoom factor, and the values on the right hand side are the mouse location on the image and the corresponding pixel value.

The tool automatically scales the image for optimum viewing, but the user may adjust the brightness and contrast using the slider bars located directly under the toolbar. The brightness and contrast may be reset to their default value using the toolbar button with the brightness and contrast icons. The default brightness and contrast values for the initial frame of the sequence will be used for every image in the sequence unless the 'Automatic Gain Control' option is checked under the 'Options' menu. If this option is checked, the optimum brightness and contrast values will be recalculated for each frame that is played. Calculating these values for every frame will slow down the playback process, but depending on the speed of the computer, it may not be noticeable.

The image may be zoomed and re-centered using the toolbar buttons with the magnifying glasses and the crosshairs respectively. All of these buttons persist until they are pushed again or until another button is pushed. This persistence allows the user to zoom in or out or re-center as many times as needed without having to push the toolbar button each time.

GROUND-TRUTHING PROCESS

Ground-truthing may be done in forward or reverse mode, which means starting either at the beginning or at the end of the sequence respectively. The ground-truthing direction may be set under the 'Options' menu as shown in *Figure 2*. Reverse mode is usually preferred since the target is easier to outline at the end of a sequence. The outline may then be adjusted as the sequence is played backward.

Once an image is opened, the target should be outlined. For reverse ground-truth mode, the target should be outlined on the last frame, or for forward ground-truth mode, the target should be outlined on the first frame. The outlining is done using the toolbar button that has an icon of a dotted outline of a rectangle. Once this button is clicked, the outline may be drawn around the target. Once drawn, the rectangle may be resized and moved using the mouse if adjustments are needed.

Since the target box will need to be adjusted between frames as the sequence is played, it may be necessary to slow down the playback speed. This can be done by adjusting the play timer, which is under the 'Options' menu, to the desired rate. The available playback rates are 10, 100, 500, 1000, and 2000 msec/frame. A good rule of thumb is to start out at the slowest rate and gradually increase the rate as needed.

Ground-truth mode must be selected for the target box to be recorded while the sequence is played. This mode may be activated by checking it under the 'Modes' menu as shown in *Figure 3*. Once this mode is turned on, the target box will be recorded every time the frame is advanced in the direction in which the ground-truth direction is set. If the ground-truth direction is reverse, the target box will only be recorded when the sequence is played in reverse and when the ground-truth direction is forward, the target box will only be recorded when the sequence is played forward.

Keystrokes should be used to advance the frames while in ground-truth mode because mouse motion moves the target box while in this mode. The up arrow on the stand-alone arrow pad which causes the



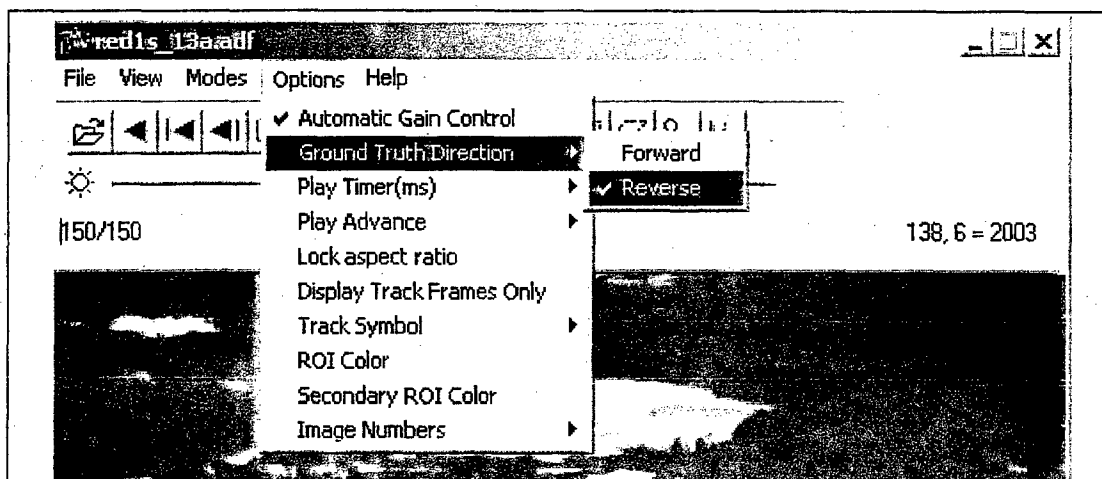


Figure 2. Options Menu




Figure 3. Modes Menu

sequence to play, will cause the sequence to play backward if reverse ground-truth mode is in effect. As the sequence plays and the target box is recorded, the box may be resized using the arrow keys on the number pad. The up and down arrows cause the height of the box to increase and decrease respectively, and the right and left arrows cause the width of the box to increase and decrease respectively. The aspect ratio of the original target box may be maintained automatically by checking the 'Lock aspect ratio' under the 'Options' menu. If this option is selected, the width will be automatically adjusted to maintain the original aspect ratio when the user adjusts the height and the height will be automatically adjusted when the user adjusts the width.

Once all the frames in the sequence have been ground-truthed, the target box information should be saved. This may be done by selecting the 'Save Ground Truth' command under the 'File' menu. Ground-truth information is usually saved with a 'gtr' file extension. The tool will choose a default filename that will be the image filename with a 'gtr' extension. This filename may be accepted, or the user may choose any filename allowable under Windows. For long sequences, it is recommended that information be saved periodically throughout the ground-truthing process.

REPLAYING A GROUND-TRUTH OR TRACKER FILE

Once a ground-truth file is created, the user may wish to view the target box information that has been saved. To do this, the file should be opened using the 'Open tracker file' under the 'File' menu. The default file extension for this operation is 'trk', but the file filter may be changed to see 'gtr' files. A tracker file may also be opened by dragging and dropping it into the image display area. However, it must have a 'trk' extension to be opened as a tracker file using this method. Next, the 'Tracker Input' mode should be checked under the 'Modes' menu. If this mode is selected and a tracker file has not been opened, the 'Open' dialog box will appear, and the user must select a file before the mode becomes active.

Once this mode is selected, a target box should appear over the image as long as the tracker or ground-truth file contained information for the current frame. As the sequence is played in this mode, the target box corresponding to each frame will be overlaid onto the image. This feature allows the user to check his or her work to make sure the target box was recorded as intended.

This mode is often used, and was designed, to view tracker simulation results. The tracker simulation output files must have a similar format to ground-truth files as described in the next section. Often tracker simulations do not process every frame, so when the file is played back, the target box appears to be blinking because it only appears on the processed frames. This effect can be bypassed by checking the 'Display Track Frames Only' item under the 'Options' menu. This option checks to make sure track information is available for a frame before playing it, and if not available, the frame is skipped.

There is also a GT (ground-truth)/Tracker Overlay mode, which may be used to compare two ground-truth and/or tracker files. The two files used in this mode are called the "tracker file" and the "GT overlay file". Both of these files may be opened using their respective commands under the 'File' menu. If either of these files is not opened before selecting the 'GT/Tracker Overlay Mode' command under the 'Modes' menu, 'Open' dialog boxes will appear to select them before this mode becomes active.

The symbol displayed for the track box may be changed to a crosshair instead of a box, or both box and crosshairs may be displayed simultaneously. This setting may be changed using the 'Track Symbol' item under the 'Options' menu.

GROUND-TRUTH AND TRACKER FILE FORMATS

Ground-truth files saved using the TAG tool use a simple tab delimited ASCII format. The first three lines are header lines containing the following information:

- Line 1 – Corresponding image file, date and time ground truth file was saved
- Line 2 – TAG version number
- Line 3 – Column headings

All remaining lines are ground-truth information in tab delimited format with the following columns: frame number, top left row coordinate, top left column coordinate, height, width, slant range, row velocity, and column velocity. The slant range is extracted from the image sequence if it is available. The row and column velocities are generated based on shift in the ground truth box's center between the previous and current frame. The velocities are used as stabilization data in some tracker simulations.

Tracker files are expected to have a very similar format. They are also tab or space delimited ASCII files. Any line that does not begin with a number, or space and then a number, will be ignored. The first five columns of the remaining lines should have the same information as the ground-truth file. These columns are frame number, top left row coordinate, top left column coordinate, height and width. Any subsequent columns on a line will be ignored. It should be noted that the ground-truth file format conforms to the tracker file format. Therefore, when using tracker input mode, a ground-truth file may be used as instructed in the previous section.

MODIFYING A GROUND-TRUTH FILE

Sometimes it may be necessary to save a ground-truth file before the ground-truthing process is complete and finish it later, or to modify an existing ground-truth file to fix an error. First, the image numbers for which modification are required should be determined. The image file should be opened and advanced to one of the frames to be modified, or if finishing a ground-truth process, it should be advanced to the next frame for which ground-truth information is required. The ground-truth file to be modified should then be opened using the 'Open ground-truth' command under the 'File' menu. Ground-truth mode should then be selected, or made active.

Once ground-truth mode is active, any advancement of the sequence in the ground-truth direction will cause the ground-truth information for the previous frame to be updated (i.e. if the ground-truth direction is set to 'reverse,' advancement in the forward direction will not record the target box). Updating occurs

whether or not ground-truth information for that frame already existed. Therefore, care should be taken so that previous work is not overwritten. The target box should be positioned at the desired location for the frame to be modified and the sequence should be advanced one frame in the ground-truth direction, or put into play if all remaining frames are to ground-truthed.

IMAGE INFORMATION DIALOGS

The image information dialog was designed to view information about the current image without cluttering the main window of the tool. This dialog may be opened by selecting 'Image Information' under the 'View' menu, shown in *Figure 4*. Initially, this dialog box will only show Irig time and slant range. More items may be viewed by selecting them under the 'View' menu of the image information dialog box as shown in *Figure 5*.

The following metrics may be viewed if a metrics file is provided: GLCM (gray-level co-occurrence matrix) target, GLCM feature, GLCM total, DeltaT, DeltaTrss, and Scrss. A metrics file usually has a 'txt' extension. It should be an ASCII text file in tab-delimited format containing: frame number, GLCM target, GLCM feature, GLCM total, DeltaT, DeltaTrss, and Scrss. A column heading line or other header lines may be present in the file. Any line that does not start with a number will be ignored. A metrics file may be opened using the 'Open metrics' command under the 'File' menu in either the image information window or the main window, or it may be dragged and dropped into the display window if it has the 'txt'

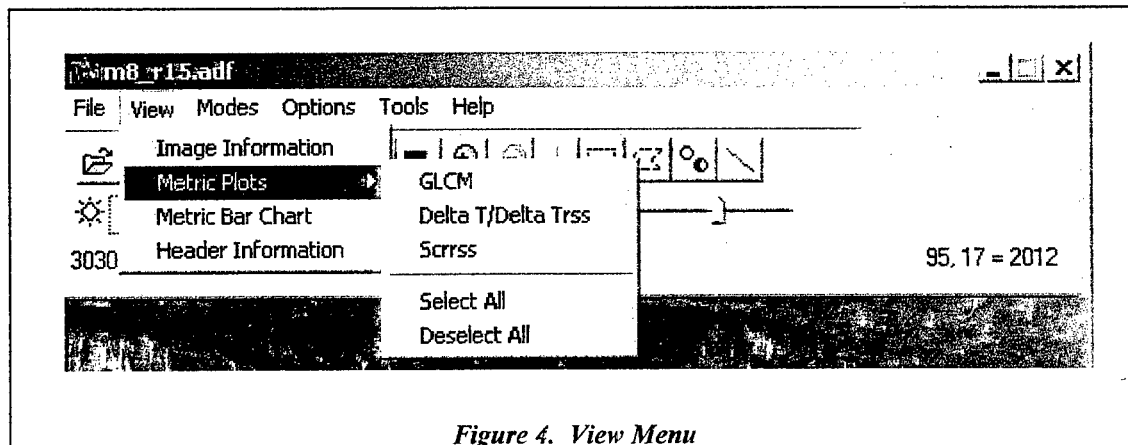

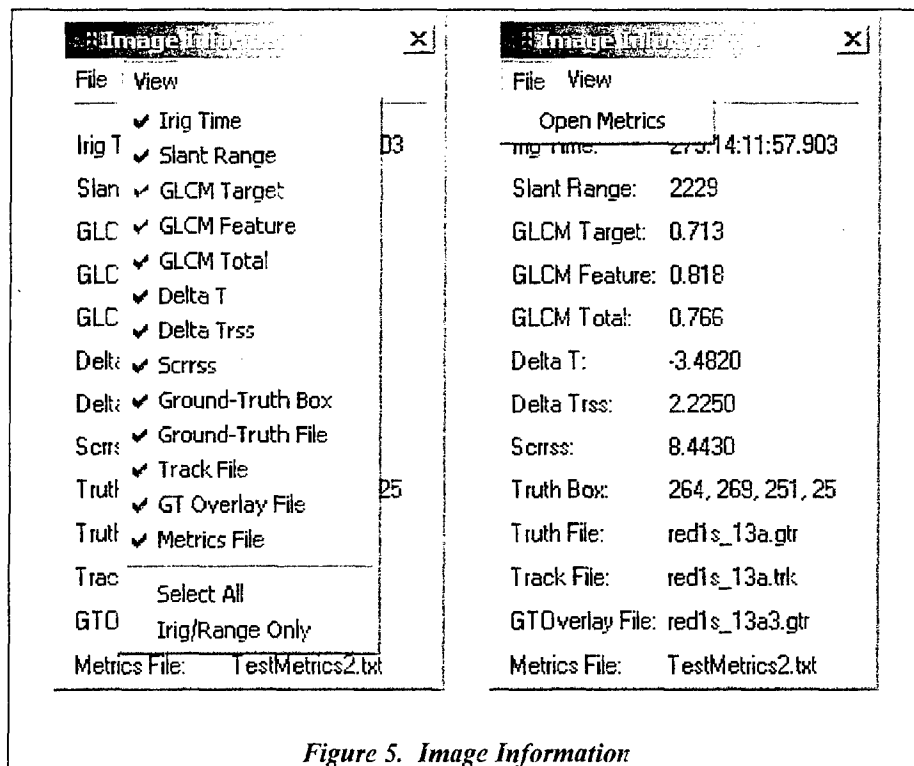



Figure 4. View Menu





extension. Other information items that may be viewed in this window are ground-truth box status, ground-truth file, tracker file, ground-truth overlay file, and metrics file. The ground-truth box status lists the coordinates of the ground-truth box for the current frame. If no ground-truth box exists for the current frame, then 'none' will be displayed. All available items may be displayed by choosing the 'Select All' command under the 'View' menu, as shown in *Figure 5*, or the default of only Irig time and slant range may be displayed by selecting 'Irig/Range Only.'

A more basic form of image information is available under the 'View' menu by using the 'Header Information' command. This shows a simple list of header information items and the corresponding value for the item on that list which is selected. This information is only available for the ADIF and NDIF file formats. *Figure 6* shows the header information dialog box.



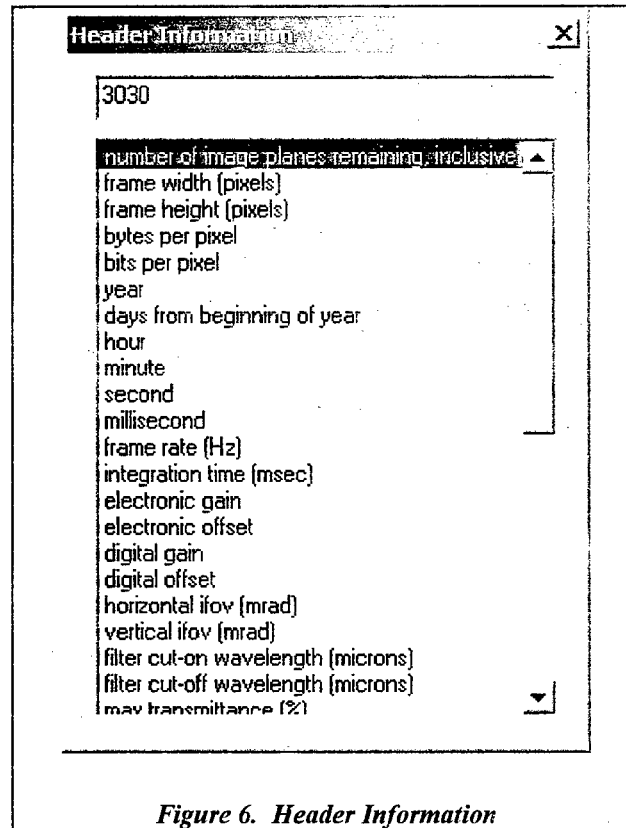


Figure 6. Header Information

METRIC PLOTS AND BAR CHART

In addition to the metric values displayed in the image information dialog box, metrics may also be displayed in plot and bar chart format. These charts may be accessed through the 'View' menu shown in *Figure 4*. A metrics file must already be opened in order for these menu items to be enabled. Each plot may be selected individually through the 'Metric Plots' submenu, or they may all be selected or deselected as a unit using the two commands at the bottom of the submenu. *Figure 7* shows the bar chart and all three metric plots.

The bar chart is a good way to compare the different metrics for a single frame, while the metric plots are a good way to see the progression of a single metric throughout a sequence's history. The metric plots are like strip charts. They are generated as the sequence is played by adding a value to the right side of the plot each time a new frame is displayed and pushing values off the left side once the plot is filled. If the sequence is stopped, the plots are not cleared, so the values displayed on the plots do not necessarily represent sequential frames in the sequence. The plots are only cleared when they are closed, or when a new sequence is opened. The DeltaT/DeltaTrss plot shows two metric values since they are very closely related. Delta T is shown in red and DeltaTrss is shown in blue.

TOOLS MENU

There are four tools available under the 'Tools' menu, which is shown in *Figure 8*: 'Histogram', 'Interpolate Slant Ranges', 'Interpolate Ground Truth', and 'Blackbody Calibration'. There is a fifth menu item, 'Record Avi', that has not yet been implemented.

A histogram of a rectangular area may be viewed by choosing the 'Histogram' tool. A rectangle must first be outlined for the histogram to be displayed. If a rectangle is not selected, an error message will appear. The limits of the histogram plot will be the minimum and maximum values of the rectangular area. More or less detail may be shown by changing the number of bins or the bin width. When one of these values is

changed, the other value is automatically calculated and changed accordingly. *Figure 9* shows an example histogram.




Figure 7. Metric Plots and Bar Chart

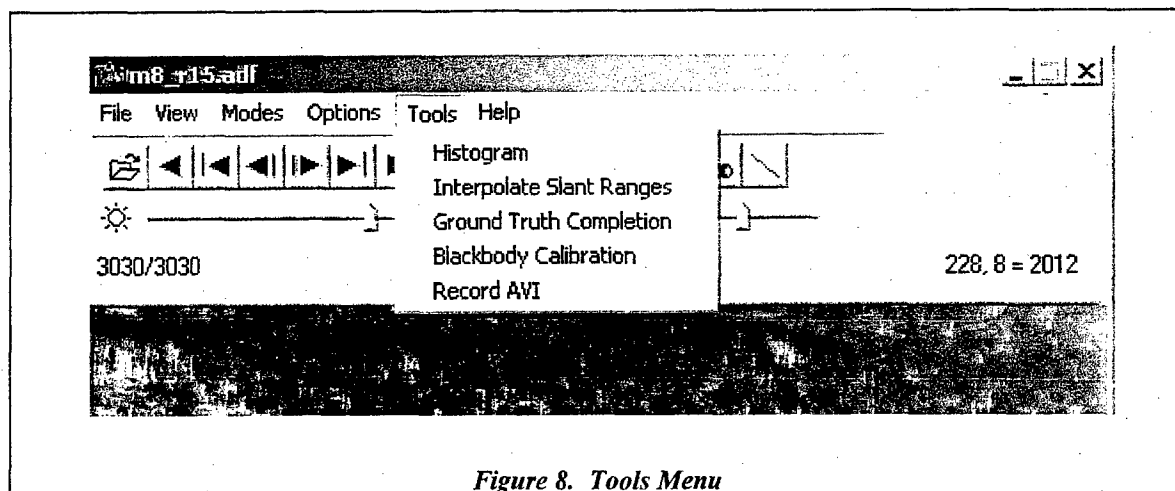


Figure 8. Tools Menu




Figure 9. Histogram Tool

The 'Interpolate Slant Ranges' tool was added in conjunction with the 'Auto-size ROI' option. In order to facilitate accurate auto-resizing of the target box, the slant range information often needs to be smoothed. For example, if slant range information is updated once per second and the frame rate of the imagery is 30 Hz, then the slant range only changes for each 30 frames of imagery, giving a stair-step effect.. The tool simply performs a moving average over the existing slant range data and inserts the new slant range data into the sequence. The user must input the number of frames over which to perform each moving average calculation. For the example just given, the number to input should be 30, as shown in *Figure 10*.

The next tool on the Tools menu is the 'Interpolate Ground-Truth' tool, shown in *Figure 11*. This tool simply sets up a command line for a separate executable (GTR_Complete.exe). When the tool is run, a separate thread is created, so that the user may continue using the software while the ground-truth completion runs. Ground-truth completion is performed using a MADD correlation, using the previous and next ground-truth rectangles. The user must specify current and new filenames, which save the current and completed ground-truth information, and a desired search area (in pixels) for the MADD correlation.

In order to use the ground-truth completion tool, the currently loaded ground-truth must be evenly spaced. This ground-truth may be loaded using the 'Open ground-truth' command under the 'File' menu, or by generating the ground-truth and immediately using the completion tool. A study was done using both easy and hard to track targets, which showed that it is optimum to skip ten frames during the ground-truthing process if using the ground-truth completion tool. Skipping frames during the ground-truthing process is done by setting the 'Play Advance' under the 'Options' menu.

The last tool on the 'Tools' menu is the blackbody calibration tool, shown in **Figure 12**. This tool will take up to ten calibration points, and will generate a 1st (linear) to a 6th order polynomial. Blackbody regions (ROI's) may be defined by outlining a region on the image with the rectangle tool and hitting the 'Get ROI' button, or may be manually entered by typing the top, bottom, left, and right coordinates of the region in a comma delimited format. Temperatures corresponding to each blackbody region should be entered in the second column. The number of blackbody points entered must exceed the order of fit desired for the resulting polynomial. Results may be saved once the calibration is performed. The output file will contain information about how the calibration coefficients were calculated, including the blackbody regions and temperatures, the image file used, and the frame used within the image sequence.




Figure 10. Interpolate Slant Ranges Tool

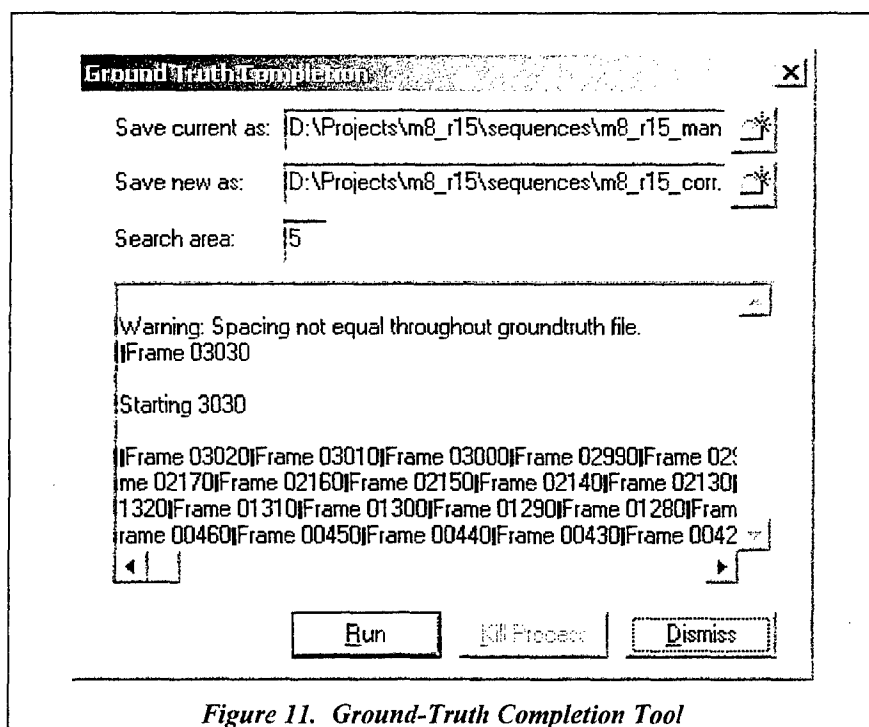


Figure 11. Ground-Truth Completion Tool




Figure 12. Blackbody Calibration Tool

ROI COLORS

The colors of the rectangular overlays, or ROIs (regions of interest) may be chosen by the user. This may be especially useful for the GT/Tracker Overlay mode to distinguish the track symbols. Users may choose the colors that they prefer at any time by selecting the 'ROI Color' or 'Secondary ROI Color' menu items under the 'Options' menu which was shown in *Figure 2*. The colors will be stored in the Windows registry, so that they do not revert to the default every time the application is started.

HELP MENU

A few help items, as well as an 'About' box are available through the 'Help' menu shown in *Figure 13*. The 'About' box, shown in *Figure 14*, displays contact information for Dynetics, Inc, and gives acknowledgement to AMCOM (Aviation and Missile Command) for their sponsorship of the development of this application. The 'Help' dialog box, shown in *Figure 15*, currently contains four help items and more are planned as development of the software continues.




Figure 13. Help Menu

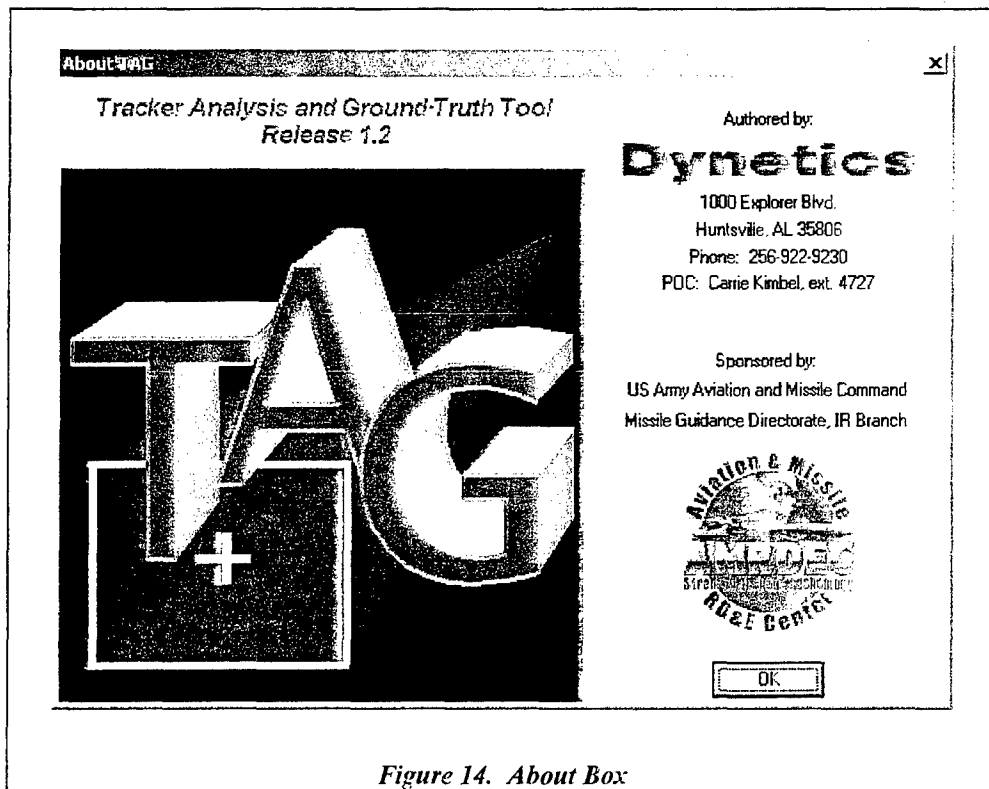


Figure 14. About Box

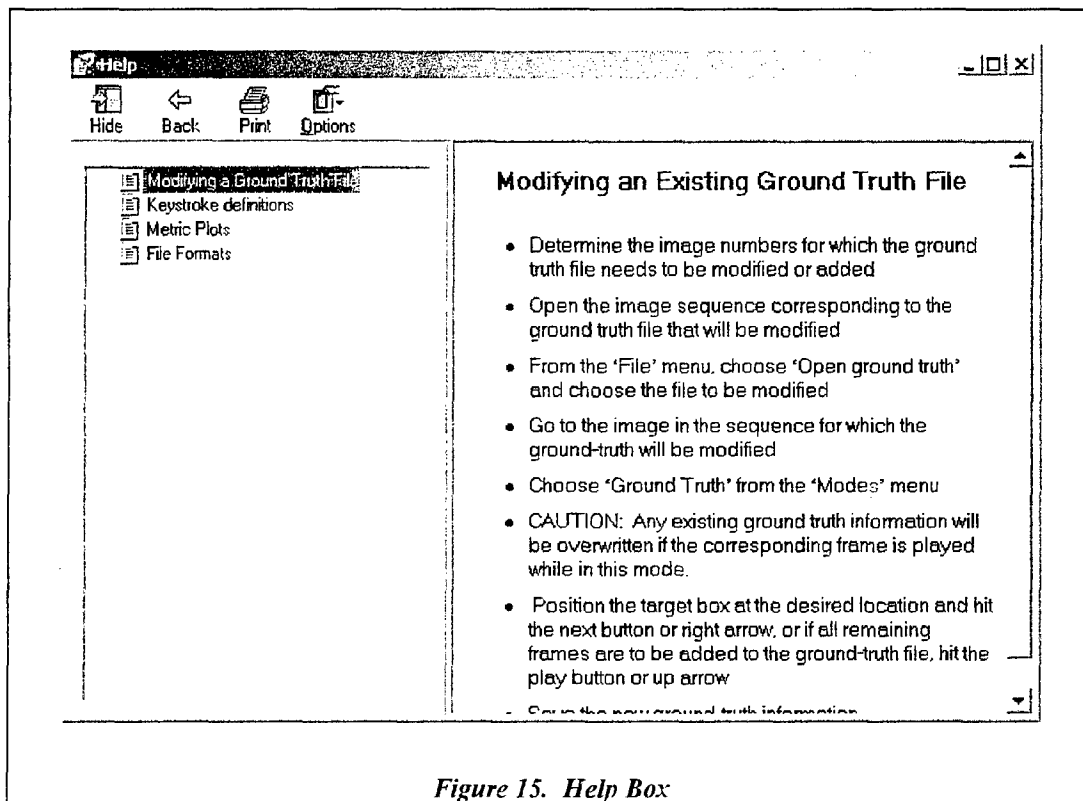


Figure 15. Help Box

CONCLUSION

The TAG tool has successfully streamlined the process of ground-truthing long image sequences for metric and tracker evaluation. During its development several features were created to make the process even easier, such as the ability to lock the aspect ratio of the ground-truth box, and the ability to ground-

truth in reverse mode. Features such as the tracker playback and ground-truth/tracker overlay make it a versatile tool that can be useful even for those not needing to generate ground-truth information. Future improvements to the tool are planned, including expanding the help information, adding a tool to create avi movies, and adding a tool to remove bad frames from a sequence.

ACKNOWLEDGEMENTS

This paper was based on work that Dynetics performed for the US AMRDEC, Missile Guidance Directorate, under the direction of Mr. Ron Passmore and Mrs. Lisa Cannon.